

With SGI's MPI and Intel OpenMP

Category: Porting to Pleiades

DRAFT

This article is being reviewed for completeness and technical accuracy.

Building Applications

To build an MPI/OpenMP hybrid executable using SGI's MPT and Intel's OpenMP libraries, your code needs to be compiled with the *-openmp* flag and linked with the *-mpi* flag.

```
%module load comp-intel/11.1.072 mpi-sgi/mpt.2.04.10789
%ifort -o your_executable prog.f -openmp -lmpi
```

Running Applications

Here is a sample PBS script for running MPI/OpenMP application on Pleiades using 3 nodes and on each node, 4 MPI processes with 2 OpenMP threads per MPI process.

```
#PBS -lselect=3:ncpus=8:mpiprocs=4:model=neh
#PBS -lwalltime=1:00:00

module load comp-intel/11.1.072 mpi-sgi/mpt.2.04.10789
setenv OMP_NUM_THREADS 2

cd $PBS_O_WORKDIR

mpiexec ./your_executable
```

You can specify the number of threads, *ompthreads*, on the PBS resource request line, which will cause the PBS prologue to set the OMP_NUM_THREADS environment variable.

```
#PBS -lselect=3:ncpus=8:mpiprocs=4:ompthreads=2:model=neh
#PBS -lwalltime=1:00:00

module load comp-intel/11.1.072 mpi-sgi/mpt.2.04.10789

cd $PBS_O_WORKDIR

mpiexec ./your_executable
```

Performance Issues

For pure MPI codes built with SGI's MPT library, performance on Nehalem-EP and Westmere-EP nodes improves by pinning the processes through setting MPI_DSM_DISTRIBUTE environment variables to 1 (or true). However, for MPI/OpenMP codes, all the OpenMP threads for the same MPI process have the same process ID and setting this variable to 1 causes all OpenMP threads to be pinned on the same core and the performance suffers.

It is recommended that MPI_DSM_DISTRIBUTE is set to 0 and *omplace* is to be used for pinning instead.

If you use Intel version 10.1.015 or later, you should also set KMP_AFFINITY to *disabled* or OMPLACE_AFFINITY_COMPAT to *ON* as Intel's thread affinity interface would interfere with dplace and oplace.

```
#PBS -lselect=3:ncpus=8:mpiprocs=4:ompthreads=2:model=neh
#PBS -lwalltime=1:00:00
```

```
module load comp-intel/11.1.072 mpi-sgi/mpt.2.04.10789
```

```
setenv MPI_DSM_DISTRIBUTE 0
setnev KMP_AFFINITY disabled
```

```
cd $PBS_O_WORKDIR
```

```
mpiexec -np 4 omplace ./your_executable
```

Article ID: 104

Last updated: 03 Oct, 2011

Computing at NAS -> Porting & Developing Applications -> Porting to Pleiades -> With SGI's MPI and Intel OpenMP

<http://www.nas.nasa.gov/hecc/support/kb/entry/104/?ajax=1>